arXiv:1406.3183v3 [stat.CO] 27 Nov 2014

Approximations of the Optimal Importance Density using Gaussian Particle Flow Importance Sampling

Pete Bunch and Simon Godsill

Abstract

Recently developed *particle flow* algorithms provide an alternative to importance sampling for drawing particles from a posterior distribution, and a number of particle filters based on this principle have been proposed. Samples are drawn from the prior and then moved according to some dynamics over an interval of pseudo-time such that their final values are distributed according to the desired posterior. In practice, implementing a particle flow sampler requires multiple layers of approximation, with the result that the final samples do not in general have the correct posterior distribution. In this paper we consider using an *approximate Gaussian flow* for sampling with a class of nonlinear Gaussian models. We use the particle flow within an importance sampler, correcting for the discrepancy between the target and actual densities with importance weights. We present a suitable numerical integration procedure for use with this flow and an accompanying step-size control algorithm. In a filtering context, we use the particle flow to sample from the optimal importance density, rather than the filtering density itself, avoiding the need to make analytical or numerical approximations of the predictive density. Simulations using particle flow importance sampling within a particle filter demonstrate significant improvement over standard approximations of the optimal importance density, and the algorithm falls within the standard sequential Monte Carlo framework.

1 Introduction

The particle filter is a Monte Carlo algorithm used for sequential inference of a filtering distribution associated with a state-space model. A set of weighted samples is advanced through time, drawn approximately from the filtering distribution. For a comprehensive introduction, see for example (Cappé et al., 2007; Doucet and Johansen, 2009). The desired posterior filtering densities contain an intractable normalising constant, which is circumvented through the use of importance sampling. The principal

challenge then, when designing a particle filter, is the selection of the importance density.

For nonlinear models, good choices of importance densities are frequently not obvious, particularly when informative observations of the latent state are made. In this situation, simple strategies such as sampling from the prior lead to a set of particles which are spread widely over the state space, of which a large proportion will have very low likelihood. The result is that the variance of the particle weights is high, and the resulting Monte Carlo estimates are dominated by a few particles with high weights. This phenomenon is known as *weight degeneracy*. Although the optimal importance density (OID) which minimises the incremental weight variance is known, it rarely has an analytical form. In practice, Gaussian approximations of the OID based on linearisation or the unscented transform are popular choices for the importance density (Doucet et al., 2000; Van Der Merwe et al., 2000), but these are not always effective.

One way in which weight degeneracy may be mitigated is by introducing the effect of each observation gradually, so that particles may be progressively drawn towards peaks in the likelihood. This can be achieved by using a discrete set of *bridging distributions* which transition smoothly between the prior and posterior. Each one is targeted in turn using importance sampling, and the accumulation of weight variance is curtailed through the use of resampling and Markov chain Monte Carlo (MCMC) steps. Such schemes have been suggested by Neal (2001); Del Moral et al. (2006) for static inference and by Godsill and Clapp (2001); Gall et al. (2007); Deutscher et al. (2000); Oudjane and Musso (2000) for particle filters.

It is possible to take the idea of bridging distributions to a limit and define a continuous sequence of distributions between the prior and the posterior. This idea was used by Gelman and Meng (1998) for the related task of simulating normalising constants, and has been used to design sophisticated assumed-density filters (Hanebeck and Feiermann, 2003; Hanebeck and Steinbring, 2012; Hagmar et al., 2011). More recently, particle filters have appeared which exploit the same principle, including the *particle flow* methods described in series of papers including (Daum and Huang, 2008, 2011), and the *optimal transport* methods of Reich (2011, 2012, 2013). A particle is first sampled from the prior (i.e. the transition) density, and then moved continuously according to some differential equation over an interval of *pseudo-time*, such that the evolution in the density corresponds to the progressive introduction of the likelihood.

Although theoretically elegant and powerful, practical implementation of optimal transport or particle flow methods require a host of approximations to be made. First, the expressions for the optimal flow dynamics are the solution to a partial differential equation and are rarely analytically tractable. Second, when applying particle flow to sample from the filtering density, the prior is generally not known analytically, and must itself be approximated. Third, once an appropriate flow has been identified, it must usually then be integrated numerically.

In this paper we focus on models which have a Gaussian prior and likelihood, but a nonlinear relationship between observations and latent states. We move the particles

according to an *approximate Gaussian flow*, based on a simple linearisation around each particle state. Unlike existing particle flow algorithms, we do not treat these directly as samples from the posterior, but as proposals in an importance sampler. Thus we obtain an accompanying differential equation for the importance weights in order to correct for the discrepancies introduced by approximating the flow. (We note that Reich (2013) has also recently suggested using a particle flow for importance sampling, but using completely different mechanisms to move the particles and update the weights.) The approximate Gaussian flow cannot be integrated analytically, so we introduce an efficient numerical scheme based on the analytical solution to the linear Gaussian flow, equipped with an effective step size control mechanism. Finally, we apply this particle flow proposal method to the OID of a particle filter, rather than to the filtering density itself. This allows the particle flow to be applied within the standard framework for particle filtering, and also avoids the need to use approximations of the predictive density.

We demonstrate the efficacy of Gaussian flow importance sampling for particle filtering with simulations on a number of challenging nonlinear models. Significant performance improvements are observed in error and effective sample size statistics.

In section 2, we review importance sampling and particle flow methods. The main exposition on using Gaussian flows for importance sampling is contained in section 3. In section 4, this strategy is applied to particle filtering, and in section 5, performance is evaluated in a number of challenging simulation studies.

A brief description of a special case of our method has been previously reported in the conference proceedings of CAMSAP (Bunch and Godsill, 2013).

2 Importance Sampling and Particle Flows

Consider the task of sampling from a Bayesian posterior distribution over a hidden state variable $x \in \mathcal{X} = \mathbb{R}^{d_x}$,

$$\pi(x) = \frac{p(x)l(x)}{K} \qquad K = \int_{\mathcal{X}} p(x)l(x)dx. \tag{1}$$

in which p and π are the prior and posterior densities respectively, which are assumed to exist, l is the likelihood and K is a normalising constant, which typically cannot be computed.

2.1 Importance Sampling

Importance sampling may be used to draw from posterior distributions (1) (Geweke, 1989; Liu, 2001). A set of N i.i.d. samples $\{x^{(i)}\}$ (or *particles*, the two terms are used interchangeably throughout) is generated according to some importance distribution with density $q(x)$ (whose support is a superset of that of $\pi(x)$) and each is assigned a

weight,

$$w^{(i)} = \frac{p(x^{(i)})l(x^{(i)})}{q(x^{(i)})} \qquad \bar{w}^{(i)} = \frac{w^{(i)}}{\sum_j w^{(j)}}. \tag{2}$$

An estimator of a posterior expectation may then be written as a finite sum over this set of weighted samples, and it is well known that this estimate is consistent, converging almost surely to its true value as the number of particles becomes large (Liu, 2001),

$$\sum_{i=1}^{\mathbf{N}} \bar{w}_n^{(i)} \phi(x^{(i)}) \overset{a.s.}{\to} \int \pi(x)\phi(x)dx. \tag{3}$$

The effectiveness of such an importance sampler depends on the choice of importance density. For integration of an arbitrary test function $\phi(x)$, it is desirable that $q(x)$ be as close to $\pi(x)$ as possible. Selecting a good importance density is therefore a foremost priority, but often proves challenging. One naive approach is to use the prior as the importance density $q(x) = p(x)$, meaning that $w^{(i)} = l(x^{(i)})$. (In a sequential setting, this is the *bootstrap filter* of Gordon et al. (1993).) This scheme is simple and easy to implement. The only requirement is that it should be possible to sample from the prior. However, it is wasteful, especially when the variance of the prior is much greater than that of the posterior, i.e. the likelihood is highly informative about the state. In this situation, the samples are widely spread over the state space, and only a few fall in the region of high likelihood. The consequence is that many have very low weight and posterior estimates are based on only a few significant particles; the resulting estimators are poor, having a high Monte Carlo variance. This is a fundamental difficulty for importance samplers. Good posterior sampling relies on having a good approximation of the posterior to begin with!

2.2 Particle Flow Sampling

Particle flow and optimal transport methods are an alternative mechanism for generating posterior samples. They have been applied to Bayesian filtering and data assimilation problems by Daum and Huang (2008, 2011, 2013); Reich (2011, 2012). The general principle is to begin with samples from the prior, then to move these according to some dynamics over an interval of pseudo-time such that the final values are distributed according to the posterior. One possible way to achieve this is to define the following geometric density sequence over the pseudo-time interval $\lambda \in [0, 1]$,

$$\pi_\lambda(x) = \frac{p(x)l(x)^\lambda}{K_\lambda} \qquad K_\lambda = \int_{\mathcal{X}} p(x)l(x)^\lambda dx. \tag{4}$$

Since $\pi_0 = p$, initial particles may be sampled from the prior. These are then moved according to an Itō stochastic differential equation (SDE) such that at every instant in

pseudo-time each one is distributed according to the appropriate density in the sequence (4),

$$dx_\lambda = \zeta_\lambda(x_\lambda)d\lambda + \eta_\lambda(x_\lambda)d\epsilon_\lambda, \tag{5}$$

in which ζ_λ and η_λ are drift and diffusion terms, and ϵ_λ is Brownian motion.

At the end, since $\pi_1 = \pi$, the final particles are independent and identically distributed according to the posterior. Hence, from the basic Monte Carlo principle, they may be used to form a consistent estimator of posterior expectations akin to (3) but with uniform weights (i.e. $w^{(i)} = 1$).

The challenge in applying such a particle flow sampler comes in finding suitable dynamics with which to move the particles such that the correct density is maintained throughout. In general, this cannot be achieved analytically, and approximations are called for (see aforesaid references). While these may sometimes lead to effective estimators, they result in the loss of consistency, and the introduction of asymptotic bias which is not easily quantified.

2.3 Exact Particle Flows

It may be shown that exact particle flows obey the following governing equation.

Theorem 2.1 *For a particle moving according to* (5)*, if the drift and diffusion are differentiable and satisfy,*

$$\begin{aligned} L(x_\lambda) - \mathbb{E}_{\pi_\lambda}[L] + \operatorname{Tr}\left[\frac{\partial \zeta_\lambda}{\partial x_\lambda}\right] + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) - \operatorname{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \Xi_\lambda}{\partial x_\lambda^2}\right] \\ - \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T D_\lambda(x_\lambda)\frac{\partial \Xi_\lambda}{\partial x_\lambda} - 2\sum_{ij}\frac{\partial D_{\lambda,ij}}{\partial x_{\lambda,i}}\frac{\partial \Xi_\lambda}{\partial x_{\lambda,j}} - \sum_{ij}\frac{\partial^2 D_{\lambda,ij}}{\partial x_{\lambda,i}\partial x_{\lambda,j}} = 0, \end{aligned} \tag{6}$$

in which

$$\begin{aligned} \Xi_\lambda(x) &= \log(\pi_\lambda(x)) & L(x) &= \log(l(x)) \\ D_\lambda(x) &= \frac{1}{2}\eta_\lambda(x)\eta_\lambda(x)^T & \mathbb{E}_{\pi_\lambda}[L] &= \int \pi_\lambda(x)L(x)dx, \end{aligned} \tag{7}$$

then the marginal density of x_λ is $\pi_\lambda(x)$ as defined by (4)*. For proof see appendix A which is based on Daum and Huang (2008).*

The governing equation relates the SDE drift and diffusion to three quantities: the gradient $\frac{\partial \Xi_\lambda}{\partial x_\lambda}$ and Hessian $\frac{\partial^2 \Xi_\lambda}{\partial x_\lambda^2}$ of the log-density at the current location, and the expected value of the log-likelihood $\mathbb{E}_{\pi_\lambda}[L]$. Intuitively, the derivative terms may be seen as controlling the particle motion due to changes in the local shape of the sequence density, while the expectation controls motion due to shifts in the bulk of the probability mass.

2.4 Particle Flow Importance Sampling

Since (6) cannot in general be solved, the approach adopted in this paper is to combine particle flow with importance sampling. Suppose that each one of a collection of particles moves according to an SDE (5), but where ζ_λ and η_λ do not satisfy (6), and the resulting density for x_λ is $\upsilon_\lambda(x) \neq \pi_\lambda(x)$. The ideal importance weight is then simply,

$$w_\lambda = \frac{\pi_\lambda(x_\lambda)}{\upsilon_\lambda(x_\lambda)}. \tag{8}$$

We can establish a differential equation for this weight.

Theorem 2.2 *A collection of particles moving according to* (5) *with differentiable drift and diffusion, and with log-weights* $W_\lambda = \log(w_\lambda)$ *evolving according to,*

$$\begin{aligned} dW_\lambda = \Bigg\{ & L(x_\lambda) - \mathbb{E}_{\pi_\lambda}\left[L\right] + \operatorname{Tr}\left[\frac{\partial \zeta_\lambda}{\partial x_\lambda}\right] + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) + \operatorname{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \Xi_\lambda}{\partial x_\lambda^2}\right] \\ & - 2\operatorname{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \Upsilon_\lambda}{\partial x_\lambda^2}\right] - \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}^T D_\lambda(x_\lambda)\frac{\partial \Upsilon_\lambda}{\partial x_\lambda} - 2\sum_{ij}\frac{\partial D_{\lambda,ij}}{\partial x_{\lambda,i}}\frac{\partial \Upsilon_\lambda}{\partial x_{\lambda,j}} \\ & - \sum_{ij}\frac{\partial^2 D_{\lambda,ij}}{\partial x_{\lambda,i}\partial x_{\lambda,j}}\Bigg\} d\lambda + \left[\frac{\partial \Xi_\lambda}{\partial x_\lambda} - \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}\right]^T \eta_\lambda(x_\lambda) d\epsilon_\lambda \end{aligned} \tag{9}$$

in which

$$\Upsilon_\lambda(x) = \log(\upsilon_\lambda(x)) \qquad \Xi_\lambda(x) = \log(\pi_\lambda(x)),$$

is properly weighted with respect to π_λ*, and the resulting weights correspond to* (8)*. For proof see appendix B*

When $\Upsilon_\lambda = \Xi_\lambda$, and ζ_λ and η_λ satisfy the conditions of theorem 2.1, then it is clear from (6) that $dW_\lambda = 0$, as we would expect when simulating perfectly from the target distribution. Furthermore, we can omit the term $\mathbb{E}_{\pi_\lambda}\left[L\right]$ in calculation since this does not depend on x_λ and thus will cancel out when the final weights are normalised. This is equivalent to using the unnormalised target density in (8).

If it were possible to simulate particle motion according to a chosen SDE, and at the same time evaluate the corresponding weights using (9), then the particles would be properly weighted importance samples for all $\lambda \in [0, 1]$, and standard convergence results would apply. In practice, it will be necessary to use approximate numerical integration schemes. Provided that these recover the ideal continuous-time evolution of both the particle state and weight as the step size tends to zero, then the resulting particle flow importance sampling will retain these asymptotic properties, but only in the limit as both the step sizes to go zero and the number of particles to infinity.

In practice, designing a numerical integration scheme which approximates (9) is not possible, because of the dependence on the unknown Υ_λ, apart from in the special case when $\eta_\lambda(x_\lambda) = 0$. Instead we show that there are other valid ways in which the weight may evolve which still result in a properly weighted collection of particles. These use the concept of targeting an extended distribution over a larger set of variables for the importance sampling.

Theorem 2.3 *A collection of particles moving according to* (5) *with differentiable drift and diffusion, and with log-weights* $W_\lambda = \log(w_\lambda)$ *evolving according to,*

$$dW_\lambda = \left\{ L(x_\lambda) - \mathbb{E}_{\pi_\lambda}[L] + \mathrm{Tr}\left[\frac{\partial \zeta_\lambda}{\partial x_\lambda}\right] + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) + \mathrm{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \Xi_\lambda}{\partial x_\lambda^2}\right] \right.$$
$$\left. - \frac{1}{2}\sum_{ijk}\left[\frac{\partial \eta_{\lambda,ik}}{\partial x_{\lambda,j}}\frac{\partial \eta_{\lambda,jk}}{\partial x_{\lambda,i}}\right] \right\} d\lambda + \sum_{ij}\frac{\partial \eta_{\lambda,ij}}{\partial x_{\lambda,i}} d\epsilon_{\lambda,j} + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \eta_\lambda d\epsilon_\lambda \quad (10)$$

is properly weighted with respect to π_λ*. The proof uses the construction of an extended target distribution which encompasses the path of the Brownian motion. See appendix C.*

Notice that (10) coincides with the ideal form (9) when $\eta_\lambda(x_\lambda) = 0$. It is possible to construct a suitable numerical scheme which results in weight evolution according to (10) as the step size tends to 0.

The idea of combining particle and importance sampling is somewhat in the spirit of (Reich, 2013), but the construction used here is substantially different, both in the type of particle flow employed and in the calculation of the weights.

3 Sampling with Gaussian Flows

3.1 Exact Gaussian Flows for Linear Gaussian Models

When the model used is linear and Gaussian, the exact flow for particle motion may be derived analytically. Suppose the likelihood takes the form of an observation $y \in \mathcal{Y} = \mathbb{R}^{d_y}$ which is linearly dependent on the state with Gaussian noise, and that the prior is also Gaussian, as follows.

Model 3.1

$$p(x) = \mathcal{N}(x\,|m_0,\, P_0) \qquad l(x) = \mathcal{N}(y\,|Hx,\, R) \quad (11)$$

P_0 *and* R *are positive definite covariance matrices.*

The following properties may be established.

Proposition 3.2 *For model 3.1, the geometric density sequence* (4) *is,*

$$\pi_\lambda(x) = \mathcal{N}\left(x \,|m_\lambda,\, P_\lambda\right) \tag{12}$$

$$P_\lambda = \left(P_0^{-1} + \lambda H^T R^{-1} H\right)^{-1} \qquad m_\lambda = P_\lambda \left[P_0^{-1} m_0 + \lambda H^T R^{-1} y\right]. \tag{13}$$

Proof The proof is straightforward using standard identities for Gaussian densities. ∎

Proposition 3.3 *A particle sampled from the prior of model 3.1 and moved according to SDE* (5) *over the interval* $[0, 1]$ *follows the density* π_λ *defined in* (12) *when the drift and diffusion terms are set to,*

$$\begin{aligned} \zeta_\lambda(x_\lambda) &= P_\lambda H^T R^{-1} \left((y - H m_\lambda) - \frac{1}{2} H (x_\lambda - m_\lambda)\right) - \frac{1}{2}\gamma (x_\lambda - m_\lambda) \\ \eta_\lambda &= \gamma^{\frac{1}{2}} P_\lambda^{\frac{1}{2}}, \end{aligned} \tag{14}$$

where $\gamma \geq 0$ *is a design parameter of the flow.*

Proof Substituting in ζ_λ and η_λ from (14), it is immediately clear that the governing equation (6) is satisfied, and hence that the flow is exact. ∎

The behaviour of the state dynamics is controlled through the choice of γ. When $\gamma = 0$, the particle motion is deterministic; when $\gamma > 0$, stochastic.

Theorem 3.4 *For model 3.1, a particle* $x_{\lambda_0} \sim \pi_{\lambda_0}$ *with this density defined in* (12) *and moved according to SDE* (5) *over the interval* $[\lambda_0, \lambda_1]$ *with drift and diffusion terms as in* (14) *reaches the state,*

$$\begin{aligned} x_{\lambda_1} = m_{\lambda_1} &+ \exp\left\{-\frac{1}{2}\gamma(\lambda_1 - \lambda_0)\right\} \left(P_{\lambda_1} P_{\lambda_0}^{-1}\right)^{\frac{1}{2}} (x_{\lambda_0} - m_{\lambda_0}) \\ &+ \left[\frac{1 - \exp\left\{-\gamma(\lambda_1 - \lambda_0)\right\}}{\lambda_1 - \lambda_0}\right]^{\frac{1}{2}} P_{\lambda_1}^{\frac{1}{2}} \left(\epsilon_{\lambda_1} - \epsilon_{\lambda_0}\right), \end{aligned} \tag{15}$$

where $A^{\frac{1}{2}}$ *is the principal matrix square root of* A*.*

Proof A constructive proof is possible by solving the SDE. This may be accomplished using a matrix integrating factor approach, and is rather lengthy. Having obtained the solution, it is straightforward to verify that it satisfies the SDE. See appendix D. ∎

Using equation (15), it is possible to calculate or sample the state at any point in pseudo-time given the state at some earlier point in pseudo-time. An example is shown in figure 1.

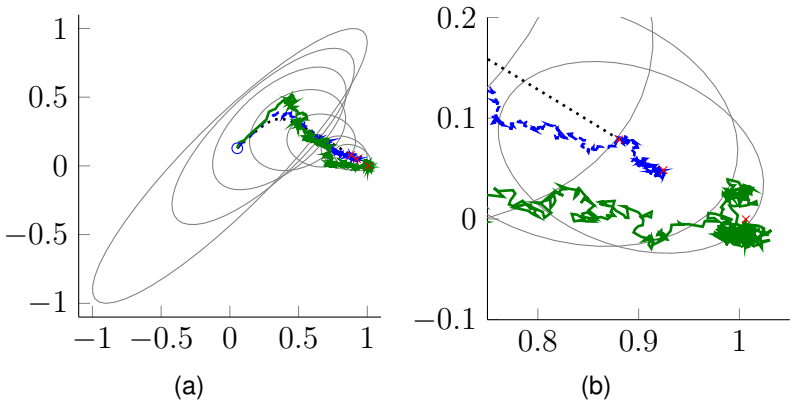

(a) (b)

Figure 1: An illustration of a Gaussian flow for a linear Gaussian model. The ellipses are one standard deviation contours of a selection of the sequence densities. The paths show the evolution of three particles from the same starting state using $\gamma = 0$ (dotted), $\gamma = 0.03$ (dashed) and $\gamma = 0.3$ (solid). The initial, prior-sampled state is shown with a circle. The second panel shows a detailed view of the final stages of the trajectories. Model parameters are $m_0 = \begin{bmatrix} 0 & 0 \end{bmatrix}^T$, $P_0 = \begin{bmatrix} 1 & 0.9 \\ 0.9 & 1 \end{bmatrix}$, $y = \begin{bmatrix} 1 & 0 \end{bmatrix}^T$, $H = I$, $R = \begin{bmatrix} 0.02 & 0.005 \\ 0.005 & 0.01 \end{bmatrix}$.

3.2 Approximate Gaussian Flows for Nonlinear Gaussian Models

For the linear Gaussian models of the previous section, sampling using a particle flow is clearly of no practical use, since the posterior distribution may be computed and sampled directly. However, it may be used as the basis of an approximately optimal flow for less tractable models. Consider the class of models with Gaussian densities but with a nonlinear dependence of the observation on the state. (N.B. In a filtering setting this encompasses the common case where both the transition and observation functions are nonlinear with additive Gaussian noise.)

Model 3.5

$$p(x) = \mathcal{N}(x\,|m_0,\ P_0) \qquad\qquad l(x) = \mathcal{N}(y\,|\psi(x),\ R) \tag{16}$$

The observation function ψ is twice differentiable with respect to x.

For nonlinear Gaussian models, the density sequence is not available analytically, nor is there a closed form expression for the particle flow. However, we can initialise the flow exactly with a sample from the Gaussian prior, and then approximate the optimal dynamics using the Gaussian flow defined in proposition 3.3. The key to this

approximation is to linearise the likelihood using a truncated Taylor expansion,

$$l(x) \approx \hat{l}(x\ ;\ x^*) = \mathcal{N}\left(\hat{y}(x^*)\,\middle|\,\hat{H}(x^*)x,\ R\right)$$

$$\hat{H}(x^*) = \left.\frac{\partial \psi}{\partial x}\right|_{x^*} \qquad \hat{y}(x^*) = y - \psi(x^*) + \hat{H}(x^*)x^*. \tag{17}$$

Using the linearised model, we can write the approximate Gaussian moments,

$$\begin{aligned}\hat{P}_\lambda(x^*) &= \left(P_0^{-1} + \lambda \hat{H}(x^*)^T R^{-1} \hat{H}(x^*)\right)^{-1} \\ \hat{m}_\lambda(x^*) &= \hat{P}_\lambda(x^*)\left[P_0^{-1} m_0 + \lambda \hat{H}(x^*)^T R^{-1} \hat{y}(x^*)\right],\end{aligned} \tag{18}$$

which parameterise the following density sequence,

$$\hat{\pi}_{\lambda|x^*}(x) = \mathcal{N}\left(x\,\middle|\,\hat{m}_\lambda(x^*),\ \hat{P}_\lambda(x^*)\right) \qquad \hat{\Xi}_{\lambda|x^*}(x) = \log\left(\hat{\pi}_{\lambda|x^*}(x)\right).$$

For this linear Gaussian approximation, the exact drift and diffusion are,

$$\begin{aligned}\hat{\zeta}_\lambda(x_\lambda\ ;\ x^*) = \hat{P}_\lambda(x^*)\hat{H}(x^*)^T R^{-1}&\left(\left(\hat{y}(x^*) - \hat{H}(x^*)\hat{m}_\lambda(x^*)\right) - \frac{1}{2}\hat{H}(x^*)(x_\lambda - \hat{m}_\lambda(x^*))\right) \\ &- \frac{1}{2}\gamma(x_\lambda - \hat{m}_\lambda(x^*)) \\ \hat{\eta}_\lambda(x^*) = \gamma^{\frac{1}{2}}\hat{P}_\lambda(x^*)^{\frac{1}{2}}&.\end{aligned} \tag{19}$$

We define the *approximate Gaussian flow* using these expressions, with the linearisation conducted around the current state,

$$dx_\lambda = \hat{\zeta}_\lambda(x_\lambda\ ;\ x_\lambda)d\lambda + \hat{\eta}_\lambda(x_\lambda)d\epsilon_\lambda \tag{20}$$

The choice of $x^* = x_\lambda$ ensures that $\hat{l}(x_\lambda\ ;\ x_\lambda) = l(x_\lambda)$, and similar equivalence for the derivatives. Using the governing equation for exact particle flows and considering both $\pi_\lambda(x)$ and $\hat{\pi}_{\lambda|x_\lambda}(x)$, it is then straightforward to show from (6) that the approximate Gaussian flow will be optimal if,

$$\mathbb{E}_{\pi_\lambda}\left[L\right] - \mathbb{E}_{\hat{\pi}_{\lambda|x_\lambda}}\left[L\right] + \text{terms involving } \frac{d^2\psi}{dx_\lambda^2} = 0. \tag{21}$$

where $\frac{d^2\psi}{dx_\lambda^2}$ is the tensor of second derivatives of the observation function.

If $\frac{d^2\psi}{dx^2} = 0$ for all x, then the model is linear and Gaussian and we recover the exact Gaussian flow. However, the flow is still optimal in the more general case where $\frac{d^2\psi}{dx_\lambda^2} = 0$ only at the current state x_λ, and also $\mathbb{E}_{\pi_\lambda}\left[L\right] - \mathbb{E}_{\hat{\pi}_{\lambda|x_\lambda}}\left[L\right] = 0$. Hence, the use

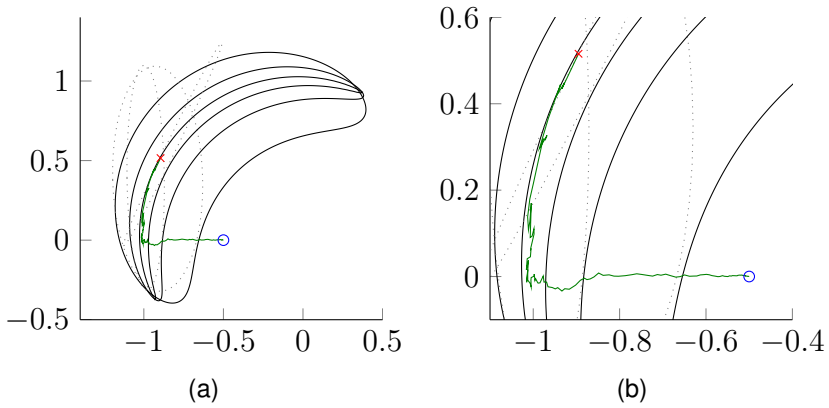

Figure 2: An illustration of an approximate Gaussian flow for a nonlinear Gaussian model. Solid contours represent the evolution of the true density sequence, and dotted contours those of the Gaussian approximations at the same times. The resulting path of a particle using $\gamma = 0.1$ is also shown. The second panel shows a close-up view of the trajectory. Model parameters are $m_0 = \begin{bmatrix} -0.4 & -0.4 \end{bmatrix}^T$, $P_0 = 0.5I$, $y = 1$, $R = 0.001$, $\psi(x) = \sqrt{x_1^2 + x_2^2}$.

of an approximate Gaussian flow implies two assumptions, that the second derivatives are small along the particle trajectory, and that the expected log-likelihood can be well-approximated using a Gaussian density.

An illustration of approximate Gaussian flow is shown in figure 2.

3.3 Numerical Integration of the Approximate Gaussian Flow

To implement the particle flow importance sampling algorithm, we need a numerical integration scheme which will allow us to approximately sample a joint trajectory for each particle state and its associated importance weight. This could be achieved using the Euler method. However, a more accurate option is available to us which exploits the analytical solution to the flow for the linear Gaussian model.

3.3.1 State Integration

The SDE for the approximate Gaussian flow can be written in the following form,

$$dx_\lambda = \left[A_\lambda(x_\lambda)x_\lambda + b_\lambda(x_\lambda)\right] d\lambda + \hat{\eta}_\lambda(x_\lambda) d\epsilon_\lambda, \tag{22}$$

where

$$A_\lambda(x_\lambda) = -\frac{1}{2}\hat{P}_\lambda(x_\lambda)\hat{H}(x_\lambda)^T R^{-1}\hat{H}(x_\lambda) - \frac{1}{2}\gamma I$$
$$b_\lambda(x_\lambda) = \hat{P}_\lambda(x_\lambda)\hat{H}(x_\lambda)^T R^{-1}\left(\hat{y}(x_\lambda) - \frac{1}{2}\hat{H}(x_\lambda)\hat{m}_\lambda(x_\lambda)\right) + \frac{1}{2}\gamma\hat{m}_\lambda(x_\lambda).$$

For an integration step from λ_0 to λ_1, the Euler method provides an approximate value of x_{λ_1} which is an exact solution to the SDE with terms fixed at λ_0 and x_{λ_0},

$$dx_\lambda = \left[A_{\lambda_0}(x_{\lambda_0})x_{\lambda_0} + b_{\lambda_0}(x_{\lambda_0})\right] d\lambda + \hat{\eta}_{\lambda_0}(x_{\lambda_0})d\epsilon_\lambda. \tag{23}$$

As the integration step size goes to 0, trajectories generated using the Euler method become exact samples according to the true SDE. For our flow, we can instead fix only the state used to form the linear approximation,

$$dx_\lambda = \left[A_\lambda(x_{\lambda_0})x_\lambda + b_\lambda(x_{\lambda_0})\right] d\lambda + \hat{\eta}_\lambda(x_{\lambda_0})d\epsilon_\lambda. \tag{24}$$

This implies matching some additional $\mathcal{O}(\delta\lambda)$ and $\mathcal{O}(\delta x)$ terms in the Taylor expansions of $\hat{\zeta}_\lambda(x_\lambda)$ and $\hat{\eta}_\lambda(x_\lambda)$. Since the resulting SDE describes a Gaussian flow, it may be solved exactly using (15), leading in this case to,

$$x_{\lambda_1} = \hat{m}_{\lambda_1}(x_{\lambda_0}) + \exp\left\{-\frac{1}{2}\gamma(\lambda_1 - \lambda_0)\right\}\left(\hat{P}_{\lambda_1}(x_{\lambda_0})\hat{P}_{\lambda_0}(x_{\lambda_0})^{-1}\right)^{\frac{1}{2}}(x_{\lambda_0} - \hat{m}_{\lambda_0}(x_{\lambda_0}))$$
$$+ \left[\frac{1 - \exp\left\{-\gamma(\lambda_1 - \lambda_0)\right\}}{\lambda_1 - \lambda_0}\right]^{\frac{1}{2}}\hat{P}_{\lambda_1}(x_{\lambda_0})^{\frac{1}{2}}(\epsilon_{\lambda_1} - \epsilon_{\lambda_0}). \tag{25}$$

From theorem 3.4, as the step size goes to 0, this recovers the ideal continuous time behaviour for the approximate Gaussian flow.

3.3.2 Weight Integration

A corresponding approximate weight update may be conducted by conditioning on the sampled value of $(\epsilon_{\lambda_1} - \epsilon_{\lambda_0})$ and using the Jacobian of (25),

$$w_{\lambda_1} = w_{\lambda_0} \times \frac{\pi_{\lambda_1}(x_{\lambda_1})}{\pi_{\lambda_0}(x_{\lambda_0})} \times \left|\frac{\partial x_{\lambda_1}}{\partial x_{\lambda_0}}\right|$$
$$\propto w_{\lambda_0} \times \frac{p(x_{\lambda_1})l(x_{\lambda_1})^{\lambda_1}}{p(x_{\lambda_0})l(x_{\lambda_0})^{\lambda_0}} \times \left|\frac{\partial x_{\lambda_1}}{\partial x_{\lambda_0}}\right|. \tag{26}$$

This is approximate, in the sense that it does not result in properly weighted samples, because the state update is not in general an invertible function due to the nonlinear dependence on x_{λ_0}. However, we can establish the following result.

Theorem 3.6 *A particle with state x_λ moved according to* (25) *and weight w_λ according to* (26)*, is properly weighted according to π_λ as the integration step size goes to 0. For proof see appendix E.*

Using the chain rule, the (i,j)th element of the Jacobian matrix is,

$$\begin{aligned}\left[\frac{\partial x_{\lambda_1}}{\partial x_{\lambda_0}}\right]_{i,j} &= \left[\frac{\partial \hat{m}_{\lambda_1}}{\partial x_{\lambda_0}} + \exp\left\{-\frac{1}{2}\gamma(\lambda_1-\lambda_0)\right\}\left(\hat{P}_{\lambda_1}(x_{\lambda_0})\hat{P}_{\lambda_0}(x_{\lambda_0})^{-1}\right)^{\frac{1}{2}}\left(I - \frac{\partial \hat{m}_{\lambda_0}}{\partial x_{\lambda_0}}\right)\right]_{i,j} \\ &+ \left(\frac{1-\exp\left\{-\gamma(\lambda_1-\lambda_0)\right\}}{\lambda_1-\lambda_0}\right)^{\frac{1}{2}}\sum_k \left[\frac{\partial \hat{P}_{\lambda_1}^{\frac{1}{2}}}{\partial x_{\lambda_0,j}}\right]_{i,k}\left[\epsilon_{\lambda_1}-\epsilon_{\lambda_0}\right]_k \\ &+ \exp\left\{-\frac{1}{2}\gamma(\lambda_1-\lambda_0)\right\}\sum_k \left[\frac{\partial\left(\hat{P}_{\lambda_1}\hat{P}_{\lambda_0}^{-1}\right)^{\frac{1}{2}}}{\partial x_{\lambda_0,j}}\right]_{i,k}(x_{\lambda_0}-\hat{m}_{\lambda_0}) \end{aligned} \tag{27}$$

The (j)th column of the derivative of $\hat{m}_\lambda(x_{\lambda_0})$ is given by,

$$\begin{aligned}\frac{\partial \hat{m}_\lambda}{\partial x_{\lambda_0,j}} = \lambda\hat{P}_\lambda(x_{\lambda_0})\Bigg(&\frac{\partial \hat{H}^T}{\partial x_{\lambda_0,j}}R^{-1}\left(\hat{y}(x_{\lambda_0}) - \hat{H}(x_{\lambda_0})\hat{m}_\lambda(x_{\lambda_0})\right) \\ &+ \hat{H}(x_{\lambda_0})^T R^{-1}\frac{\partial \hat{H}}{\partial x_{\lambda_0,j}}\left(x_{\lambda_0} - \hat{m}_\lambda(x_{\lambda_0})\right)\Bigg),\end{aligned} \tag{28}$$

where $\frac{\partial \hat{H}^T}{\partial x_{\lambda_0,j}}$ is a matrix whose (i,k)th term is $\frac{\partial^2 \psi_i}{\partial x_{\lambda_0,j}\partial x_{\lambda_0,k}}$. The two matrix square root derivatives may be evaluated by observing that since $A^{\frac{1}{2}}A^{\frac{1}{2}} = A$, then,

$$A^{\frac{1}{2}}\frac{\partial A^{\frac{1}{2}}}{\partial x_{\lambda_0,j}} + \frac{\partial A^{\frac{1}{2}}}{\partial x_{\lambda_0,j}}A^{\frac{1}{2}} = \frac{\partial A}{\partial x_{\lambda_0,j}},$$

and hence the derivative may be found by solving a Sylvester equation by standard methods (Bartels and Stewart, 1972). On the right hand side of these equations we need,

$$\begin{aligned}\frac{\partial \hat{P}_{\lambda_1}}{\partial x_{\lambda_0,j}} &= -\lambda_1\hat{P}_{\lambda_1}(x_{\lambda_0})\left(\frac{\partial \hat{H}^T}{\partial x_{\lambda_0,j}}R^{-1}\hat{H}(x_{\lambda_0}) + \hat{H}(x_{\lambda_0})R^{-1}\frac{\partial \hat{H}^T}{\partial x_{\lambda_0,j}}\right)\hat{P}_{\lambda_1}(x_{\lambda_0}) \\ \frac{\partial\left(\hat{P}_{\lambda_1}\hat{P}_{\lambda_0}^{-1}\right)}{\partial x_{\lambda_0,j}} &= \hat{P}_{\lambda_1}(x_{\lambda_0})\left(\frac{\partial \hat{H}^T}{\partial x_{\lambda_0,j}}R^{-1}\hat{H}(x_{\lambda_0}) + \hat{H}(x_{\lambda_0})R^{-1}\frac{\partial \hat{H}^T}{\partial x_{\lambda_0,j}}\right) \\ &\quad\times\left(\lambda_0 I - \lambda_1\hat{P}_{\lambda_1}(x_{\lambda_0})\hat{P}_{\lambda_0}(x_{\lambda_0})^{-1}\right).\end{aligned}$$

3.3.3 Step Size Control

Effective numerical integration requires careful consideration of the integration step sizes. Using smaller step sizes reduces the associated errors, resulting in a path which more accurately represents a sample from the approximate Gaussian flow. However, in practice the number of steps needs to be kept fairly low, to minimise the computational burden. In some instances, it may be sufficient to use a fixed step size, or a predetermined time grid chosen with a tuning run. However, an adaptive scheme is preferable for greatest efficiency.

Adaptation may be conducted by forming an estimate of the error introduced by the numerical integration scheme, and adjusting the step size to keep this below a predetermined threshold. When $\gamma > 0$ and the flow is stochastic, the error is calculated conditional on the sampled path of the Brownian motion. Step size adaptation is performed independently for each particle.

To form a local error estimate for each integration step, we can compare the SDE for the approximate Gaussian flow (22) with that implied by the numerical integration (24). The difference between them describes the evolution of the integration error,

$$de_\lambda = \left[\left(A_\lambda(x_\lambda) - A_\lambda(x_{\lambda_0})\right) x_\lambda + \left(b_\lambda(x_\lambda) - b_\lambda(x_{\lambda_0})\right)\right] + \left[\hat{\eta}_\lambda(x_\lambda) - \hat{\eta}_\lambda(x_{\lambda_0})\right] d\epsilon_\lambda.$$

Integrating from λ_0 to λ_1 and approximating each integrand with the average of its initial and final value,

$$\begin{aligned} e_{\lambda_1|\lambda_0} &= \int_{\lambda_0}^{\lambda_1} \left[\left(A_\lambda(x_\lambda) - A_\lambda(x_{\lambda_0})\right) x_\lambda + \left(b_\lambda(x_\lambda) - b_\lambda(x_{\lambda_0})\right)\right] d\lambda \\ &\qquad + \int_{\lambda_0}^{\lambda_1} \left[\hat{\eta}_\lambda(x_\lambda) - \hat{\eta}_\lambda(x_{\lambda_0})\right] d\epsilon_\lambda \\ &\approx \frac{1}{2}(\lambda_1 - \lambda_0)\left[\left(A_{\lambda_1}(x_{\lambda_1}) - A_{\lambda_1}(x_{\lambda_0})\right) x_{\lambda_1} + \left(b_{\lambda_1}(x_{\lambda_1}) - b_{\lambda_0}(x_{\lambda_0})\right)\right] \\ &\qquad + \frac{1}{2}\left[\hat{\eta}_{\lambda_1}(x_{\lambda_1}) - \hat{\eta}_{\lambda_1}(x_{\lambda_0})\right](\epsilon_{\lambda_1} - \epsilon_{\lambda_0}). \end{aligned} \tag{29}$$

Pseudo-time step sizes may now be adjusted so that the magnitude of this error statistic is kept below a threshold. If the error is too large then the new state is rejected and the step size reduced. Note than when $\gamma > 0$ we must be particularly careful with this procedure. Since the weight and step size error calculations are conditioned on the sampled path of the Brownian motion ϵ_λ, we cannot simply discard this and sample afresh. Intermediate values must be drawn from a Brownian bridge conditional on the existing skeleton of sampled points.

Mechanisms for adjusting the step sizes may be borrowed directly from well-established numerical integration algorithms for solving differential equations (see for example (Shampine and Reichelt, 1997)).

Since $(\epsilon_{\lambda_1} - \epsilon_{\lambda_0})$ is of the order $\mathcal{O}\left((\lambda_1 - \lambda_0)^{\frac{1}{2}}\right)$, it is expected that the stochas-

tic term will dominate the integration error unless γ is small. This suggests that we should ordinarily use $\gamma = 0$. The advantage of using other values is in their use for implementing efficient MCMC kernels, as we discuss in section 3.6.

3.4 Summary

Particle flow importance sampling may be conducted by first drawing a set of particles from the prior, and then allowing their states and weights to evolve according to an approximate Gaussian flow over an interval of pseudo time $\lambda \in [0, 1]$ using numerical integration. Pseudo-code for the procedure is provided in algorithm 1.

Algorithm 1 Gaussian flow importance sampling for the nonlinear Gaussian model.

Require: Parameters m_0, P_0, R, y, ψ, according to model 3.5.
for $i = 1$ **to** $\mathbf{N}$ **do**
  Sample $x_0^{(i)} \sim \mathcal{N}(x \,|m_0,\, P_0)$.
  Initialise unnormalised weight $w_0^{(i)} = 1$.
  Initialise Brownian motion $\epsilon_0^{(i)} = 0$.
  Set $\lambda = 0$
  while $\lambda \leq 1$ **do**
    Increment $\lambda \leftarrow \min\{1,\, \lambda + \Delta\lambda\}$ with $\Delta\lambda$ specified by a fixed grid or adaptive method. (See section 3.3.3).
    Linearise observation function using (17).
    Sample a value for the Brownian motion ϵ_λ.
    Calculate approximate Gaussian moments using (18).
    Advance state using (25), yielding $x_\lambda^{(i)}$.
    Advance weight using (26), yielding $w_\lambda^{(i)}$.
  end while
  Set $x^{(i)} \leftarrow x_1^{(i)}$ and $w^{(i)} \leftarrow w_1^{(i)}$.
end for
Normalise weights $\bar{w}^{(i)} = w^{(i)} / \sum_j w^{(j)}$.
return Importance weighted posterior samples $\left\{x^{(i)},\, \bar{w}^{(i)}\right\}$.

3.5 Performance Characterisation

We now illustrate the the operation of particle flow importance sampling on a simple example model, and use this to explore its dependence on various model and algorithm parameters. The model is defined by the following parameters,

$$m_0 = \mathbf{1} \qquad P_0 = \sigma_x^2 I \qquad \psi(x) = \left(\sum_i x_i^2\right)^{\frac{1}{2}} \qquad R = \sigma_y^2. \tag{30}$$

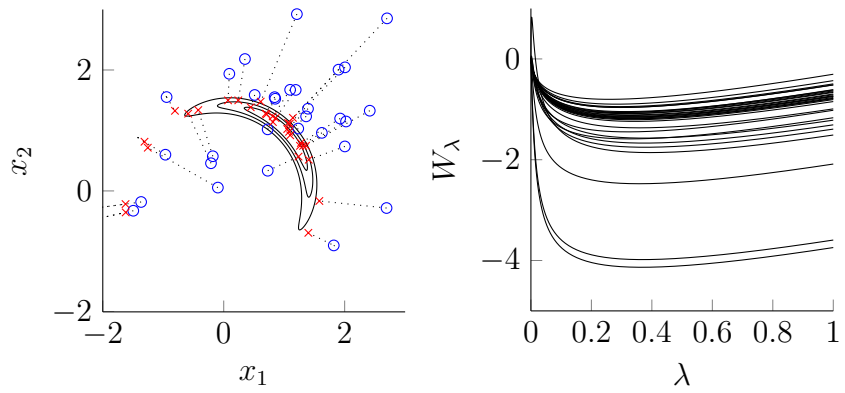

Figure 3: A simple example of Gaussian flow importance sampling using (30) with $\sigma_x = 1$ and $\sigma_y = 0.1$, showing evolution of (a) the particle states and (b) the particle log-weights. Contours of the target posterior are shown with solid lines. Particle paths are shown with dotted lines, with the initial and final state shown by a circle and cross respectively.

Figure 3 shows typical evolution of particle states and weights using a deterministic flow.

In figures 4 to 7, we compare the root mean square error (RMSE) and effective sample size (ESS) (Kong et al., 1994) obtained using particle flow importance sampling against those obtained from conventional importance sampling using two different choices of importance distribution: the prior, and a Laplace approximation of the posterior formed at the mode. For these results, the particle flow uses a fine fixed grid of pseudo-time steps, in order to give an idea of optimal performance. The dependence on the prior and observation variances and state dimension d_x are illustrated, and the effect of the diffusion scale factor γ.

The Gaussian flow consistently outperforms the simpler samplers, particularly so in the more extreme parameter settings. Particle flow sampling has the greatest advantage when the posterior is ill-conditioned or particularly non-Gaussian. This occurs when the state dimension is greater than the number of observations, and the observations are informative compared to the prior, either because the prior variance is high or the observation variance low. In combination with the nonlinear observation function, this gives rise to complex posterior distributions, in which the modes are often highly "curved" (as in figure 3) and thus poorly represented by a single Gaussian.

In figure 8 we provide another demonstration of the benefits of using particle flow importance sampling, this time using a more practical model (a single-frame of the altitude-assisted tracking described in section 5) and with an implementation of the adaptive step size mechanism. Particle states are shown before and after a resampling step. The number of samples drawn in each case is scaled such that the running time

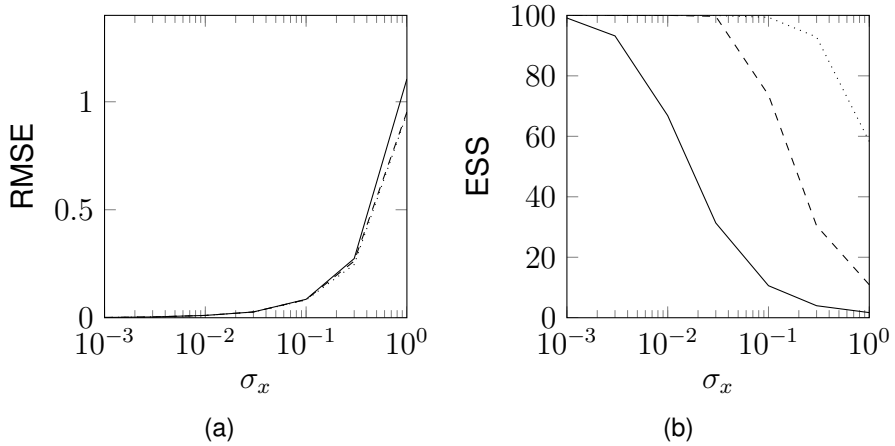

Figure 4: Root mean square error (a) and effective sample size (b) with $d_x = 2$, $\sigma_y = 0.1$ and varying prior standard deviation σ_x for three importance samplers. Sampling from the prior (solid), Laplace approximation importance density (dashed), and particle flow with $\gamma = 0$ (dotted). 100 particles for each.

for each is the same. It is clear that the particle flow is better able to characterise the posterior, while doing more than just place particles around a mode.

3.6 Resample-Move with Particle Flow Proposals

If an importance sampler generates a set of particles which is dominated by a small number with large weights, then the resulting posterior estimates will have a high variance. When this happens, a post-processing stage known as resample-move (Gilks and Berzuini, 2001) may improve the situation. The weighted particle set is first resampled according to the normalised importance weights to produce an unweighted set. In this standard procedure, low-weight particles are discarded and high-weight particles copied to replace them, with the number of replicates chosen randomly in an appropriate manner so as to ensure unbiasedness (Hol et al., 2006; Douc et al., 2005). These replicated particles are then perturbed by sampling from an MCMC kernel so as to spread them around and further explore the promising areas of the state space. Resampling reduces the weight variance of a particle set at the cost of introducing dependence between the particles. The MCMC steps are used to reduce this dependence. Note that the MCMC does not need to be run to convergence in resample-move, since it is being used merely to improve sample diversity.

Implementing resample-move effectively requires some additional algorithm parameters to be selected, such as the number of MCMC steps and an appropriate proposal distribution for Metropolis-Hastings (MH). When particle flow sampling is used, there is an obvious choice for this proposal; simply return to the original state for each particle which was sampled from the prior, x_0, and re-simulate a new path through pseudo-time.

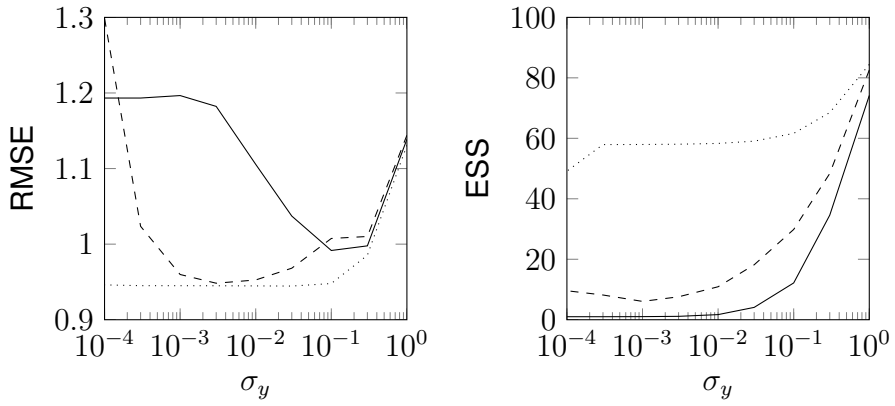

Figure 5: Root mean square error (a) and effective sample size (b) with $d_x = 2$, $\sigma_x = 1$ and varying observation standard deviation σ_y for three importance samplers. Sampling from the prior (solid), Laplace approximation importance density (dashed), and particle flow with $\gamma = 0$ (dotted). 100 particles for each.

The choice of proposal distribution is thus reduced to setting a value of γ, the diffusion scale factor, which will control the size of the proposed moves. Clearly with $\gamma = 0$ the motion is deterministic and no move would be taken, i.e. the chain remains stuck in its current location.

Since new values of x_1 are to be drawn independently conditional on x_0, the acceptance probability is simply that for an MH independence sampler. That is, if the existing state has weight w (unnormalised, before resampling), and new state for the MH proposal has unnormalised weight w^*, then the MH acceptance probability is,

$$\min\left\{1, \frac{w^*}{w}\right\}. \tag{31}$$

Figure 9 shows two stochastic flows being used for resample-move, illustrating the scope for exploring the state space using this method.

4 Applications in Particle Filtering

Our motivating purpose for studying particle flows is for use in filtering. We consider a standard discrete-time Markovian state space model in which the transition, observation and prior models have closed-form densities,

$$x_n \sim f(x_n|x_{n-1}) \qquad y_n \sim g(y_n|x_n) \qquad x_1 \sim p(x_1),$$

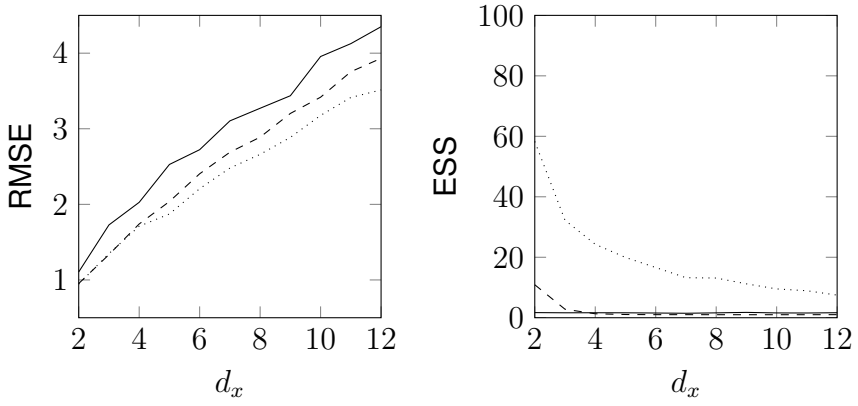

Figure 6: Root mean square error (a) and effective sample size (b) with $\sigma_x = 1$, $\sigma_y = 0.1$ and varying state dimension for three importance samplers. Sampling from the prior (solid), Laplace approximation importance density (dashed), and particle flow with $\gamma = 0$ (dotted). 100 particles for each.

where the random variable x_n is the hidden state of a system at time n, and y_n is an incomplete, noisy observation.

A conventional particle filter (Cappé et al., 2007; Doucet and Johansen, 2009) uses importance sampling to estimate distributions recursively over the path of the state variables, $x_{1:n} = \{x_1, \ldots, x_n\}$, such that,

$$\sum_{i=1}^{\mathbf{N}} \bar{w}_n^{(i)} \phi(x_{1:n}^{(i)}) \overset{a.s.}{\to} \int \pi(x_{1:n}) \phi(x_{1:n}) dx_{1:n}.$$

Each step begins by selecting a set of ancestors $\{a_n^{(i)}\}$ from amongst the $(n-1)$th step particles according to the corresponding weights. Next, a new state is proposed for each particle from an importance density $x_n^{(i)} \sim q(x_n | x_{n-1}^{(a_n^{(i)})}, y_n)$, and this is concatenated to the ancestral path to form the new particle $x_{1:n}^{(i)} \leftarrow \left\{ x_{1:n-1}^{(a_n^{(i)})}, x_n^{(i)} \right\}$. An importance weight is then assigned to the particle to account for the discrepancy between importance and target distributions,

$$w_n^{(i)} = \frac{p(x_{1:n}^{(i)} | y_{1:n})}{p(x_{1:n-1}^{(a_n^{(i)})} | y_{1:n-1}) q(x_n^{(i)} | x_{n-1}^{(a_n^{(i)})}, y_n)} \propto \frac{f(x_n^{(i)} | x_{n-1}^{(a_n^{(i)})}) g(y_n | x_n^{(i)})}{q(x_n^{(i)} | x_{n-1}^{(a_n^{(i)})}, y_n)}. \tag{32}$$

It was shown by Doucet et al. (2000) that the incremental weight variance is minimised by proposing from the conditional posterior $q(x_n | x_{n-1}^{(a_n^{(i)})}, y_n) = p(x_n | x_{n-1}^{(a_n^{(i)})}, y_n)$, known as the optimal importance density (OID). This cannot be used routinely due to

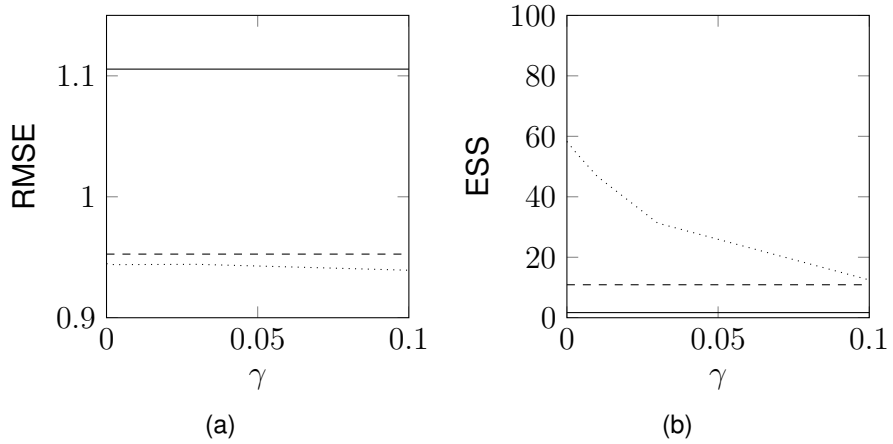

Figure 7: Root mean square error (a) and effective sample size (b) with $d_x = 1$, $\sigma_x = 1$, $\sigma_y = 0.1$ for three importance samplers. Sampling from the prior (solid), Laplace approximation importance density (dashed), and particle flow (dotted) with varying diffusion scale factor γ. 100 particles for each.

an intractable normalising constant required in the weight caluclations.

4.1 Existing Particle Flow Approaches

The approach taken by Daum and Huang (2008, 2011, 2013); Reich (2011, 2012) is to apply particle flow sampling directly to the filtering density. Assume that a set of unweighted particles exists approximating $p(x_{n-1}|y_{1:n-1})$. The predictive density at the next step is related by,

$$p(x_n|y_{1:n-1}) = \int f(x_n|x_{n-1})p(x_{n-1}|y_{1:n-1})dx_{n-1}, \tag{33}$$

which can thus be sampled by simply drawing $x_n^{(i)} \sim f(x_n|x_{n-1}^{(i)})$ for each particle and then marginalising (i.e. discarding) the old states. Defining this predictive density as the prior and the filtering density as the posterior, a particle flow is used to sample from,

$$p(x_n|y_{1:n}) = \frac{p(x_n|y_{1:n-1})g(y_n|x_n)}{K_n}. \tag{34}$$

The difficulty with this approach is that finding an appropriate flow generally requires at least the prior and often also its gradient and Hessian to be calculable pointwise. This is not the case for the predictive density, $p(x_n|y_{1:n-1})$. (Note that we could use a Monte Carlo approximation of this density, but the resulting algorithm has a complexity of $\mathcal{O}(\mathbf{N}^2)$ in the number of particles.) Reich (2011, 2012, 2013) address this by making analytical approximations of this density as a Gaussian or Gaussian mixture. Daum

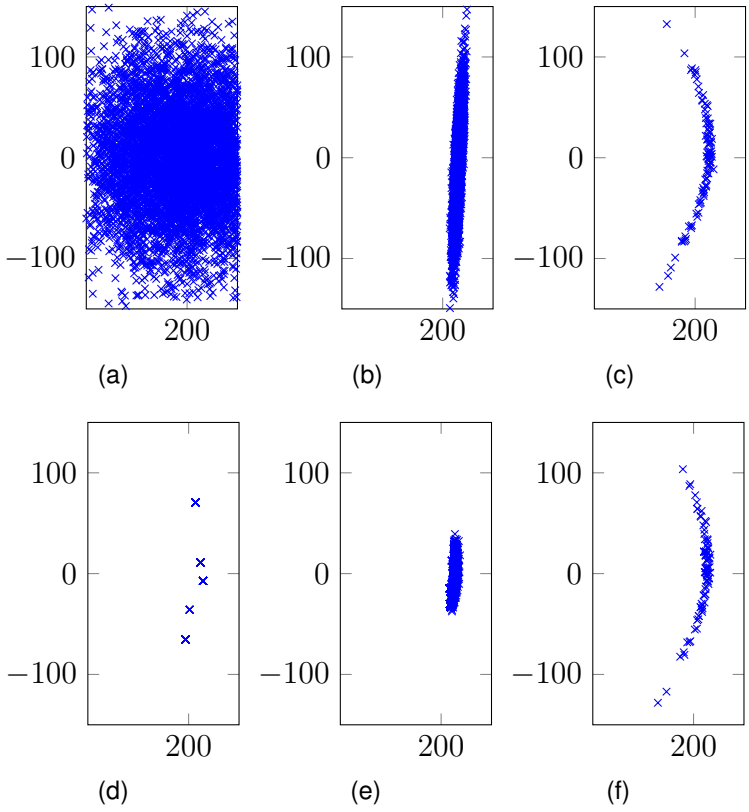

Figure 8: Particles sampled before (a),(b),(c) and after (d),(e),(f) resampling using three different strategies: sampling from the prior (a),(d), sampling from a Laplace approximation of the posterior (b),(e) and a Gaussian flow (c),(f).

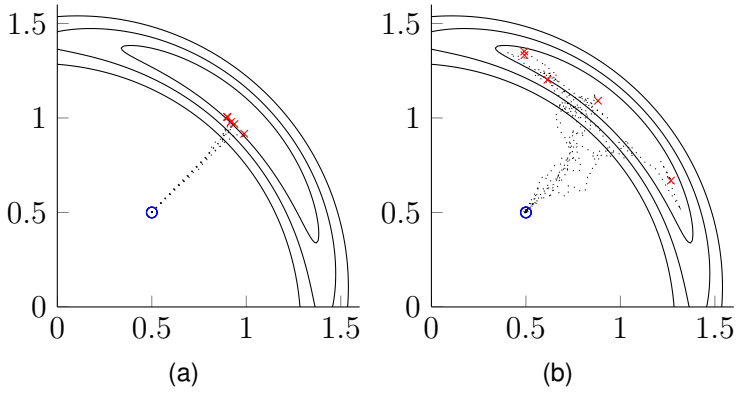

Figure 9: Particle trajectories for two particle flows targeting the example model (30) using $\sigma_x = 1$, $\sigma_y = 0.1$, using (a) $\gamma = 0.001$ and (b) $\gamma = 0.1$. Contours of the target posterior are shown with solid lines. Particle paths are shown with dotted lines, with the initial and final state shown by a circle and cross respectively. For this easy problem, the acceptance probabilities in both cases are close to 1.

and Huang (2008, 2011, 2013); Daum et al. (2009) use a number of methods, including Gaussian and various numerical approximations. These approximations alter the actual distribution of the particles. The filter no longer returns a properly weighted set of particles representing the posterior and consistent estimates of posterior expectations are no longer guaranteed.

Furthermore, the existing particle flow algorithms do not fall within the framework of ordinary particle filters. They only provide us with an estimate of the marginal filtering density $p(x_n|y_{1:n})$, rather than the more conventional path filtering density $p(x_{1:n}|y_{1:n})$. This may sometimes be all that is needed, but on other occasions samples of the entire path are essential, for example for smoothing (Kitagawa, 1996) or parameter estimation schemes, such as particle MCMC (Andrieu et al., 2010).

4.2 Gaussian Flow Approximations to the Optimal Importance Density

In this work, we use particle flow sampling within the standard particle filtering framework, thus retaining samples of the entire path and avoiding the need for additional layers of approximation. In order to achieve this, we need to consider two different density sequences. The flow for each particle state is derived by targeting the optimal importance density (OID) with the sequence,

$$\frac{f(x_n|x_{n-1}^{(a_n^{(i)})})g(y_n|x_n)^\lambda}{K_\lambda(x_{n-1}^{(a_n^{(i)})})}. \tag{35}$$

This allows us to sample a value for x_n conditional on the history $x_{1:n-1}$. Meanwhile, the weight updates are conducted so as to target the filtering density over the entire trajectory, with the sequence,

$$\frac{p(x_{1:n-1}|y_{1:n-1})f(x_n|x_{n-1}^{(a_n^{(i)})})g(y_n|x_n)^\lambda}{K_\lambda}. \tag{36}$$

With this simple modification, the required weight update formula becomes,

$$w_{\lambda_1} \propto w_{\lambda_0} \times \frac{g(y_n|x_{\lambda_1})^{\lambda_1} f(x_{\lambda_1}|x_{n-1})}{g(y_n|x_{\lambda_0})^{\lambda_0} f(x_{\lambda_0}|x_{n-1})} \times \left|\frac{\partial x_{\lambda_1}}{\partial x_{\lambda_0}}\right|. \tag{37}$$

5 Simulations

Numerical testing using simulated data is presented to demonstrate the efficacy of Gaussian flow sampling for particle filtering. We measure performance by considering RMSE values, using the empirical particle mean as a point estimate, and average effective sample size (ESS), measured before resampling (Kong et al., 1994).

The following particle filters (and their respective importance densities) were tested:

- A bootstrap filter (BF), using the transition density. (Gordon et al., 1993)
- An extended particle filter (EPF), using a Gaussian density chosen by linearisation about the predictive mean, in the style of an extended Kalman filter. (Doucet et al., 2000)
- An unscented particle filter (UPF), using a Gaussian density chosen using the unscented transform, in the style of an unscented Kalman filter. (Van Der Merwe et al., 2000)
- A Laplace approximation particle filter (LAPF), using a Gaussian density chosen by truncation of the Taylor series of the log of the unnormalised OID around a local maximum (Doucet et al., 2000). Gradient ascent is used to locate the maximum.
- A Gaussian flow particle filter (GFPF), using the the Gaussian flow importance sampling method, with $\gamma = 0$. The adaptive step size mechanism is used and requires roughly 5 to 40 steps.

The posterior filtering distributions of the chosen models can assume complex and irregular shapes, sometimes leading to the complete failure of the EPF and UPF. The LAPF is generally slow because the maximisation procedure struggles with the irregular mode shapes.

The number of particles for the GFPF was set to 100. For the remaining filters, the number of particles was increased so as to achieve a similar running time. On the altitude-assisted tracking model, the LAPF in fact took roughly 3 times as long as the other algorithms.

5.1 Models

5.1.1 Altitude-Assisted Tracking

We consider tracking a small aircraft over a mapped landscape, a scenario inspired by Schön et al. (2005). Time of flight and Doppler measurements from a radio transmitter on the aircraft provide accurate measurements of range r_n, and range rate s_n, but only a low resolution measurement of bearing b_n. In addition, accurate measurements are made of the height above the ground h_n. The profile of the terrain (i.e. the height of the ground above a datum at each point) has been mapped.

At time step n, the latent state for our model is,

$$x_n = \begin{bmatrix} p_n^T & v_n^T \end{bmatrix}^T,$$

where p_n and v_n are the 3-dimensional position and velocity of the aircraft respectively, and the observation is,

$$y_n = \begin{bmatrix} b_n & r_n & h_n & s_n \end{bmatrix}^T. \tag{38}$$

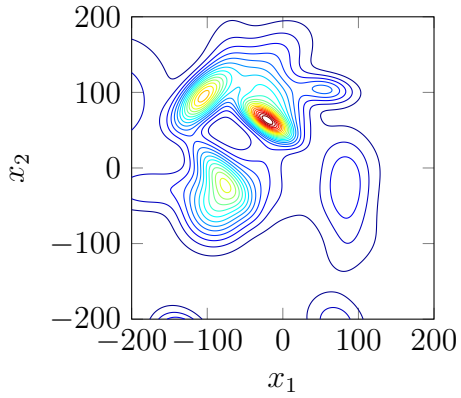

Figure 10: Contour plot of an example simulated terrain map.

The observation function is described by the following equations,

$$b_n = \arctan\left(\frac{p_{n,1}}{p_{n,2}}\right) + e_{n,1} \qquad r_n = \sqrt{p_{n,1}^2 + p_{n,3}^2 + p_{n,3}^2} + e_{n,2}$$
$$h_n = p_{n,3} - T(p_{n,1}, p_{n,2}) + e_{n,3} \qquad s_n = \frac{p_n \cdot v_n}{r_n} + e_{n,4},$$

where $T(p_{n,1}, p_{n,2})$ is the terrain height at the corresponding horizontal coordinates. The four noise terms have independent zero-mean Gaussian densities and the respective variances are $\left(\frac{\pi}{9}\right)^2$, 0.1^2, 0.1^2, 0.1^2. A linear Gaussian near-constant velocity transition model is used (Bar-Shalom and Li, 1995), with volatility of 30^2. The terrain profile was modelled as a mixture of randomly-generated Gaussian blobs. An example is shown in figure 10.

The accurate measurements of range, range rate and height constrain the region of high posterior probability to lie on a 3 dimensional subspace, which can take some very irregular shapes.

5.1.2 Fitting A Skeletal Model

We consider a toy motion-capture problem, in which camera measurements are used to estimate the pose of a human arm. The latent state consists the 3D coordinates of the shoulder joint $r_{S,}$, the orientation α_B, the angles of the shoulder α_S and elbow α_E, and the lengths of the upper d_U and lower d_L arm. The evolution of each of these is modelled as a random walk with Gaussian noise. The variances are 0.5^2 for position in the transverse directions, and 0.1^2 in the depth direction, $\frac{\pi}{18}^2$ for the angles and 0.001^2 for the lengths (which allows for model inaccuracy, and avoids the need to do static parameter estimation). The observation model consists of two stages. First, the elbow

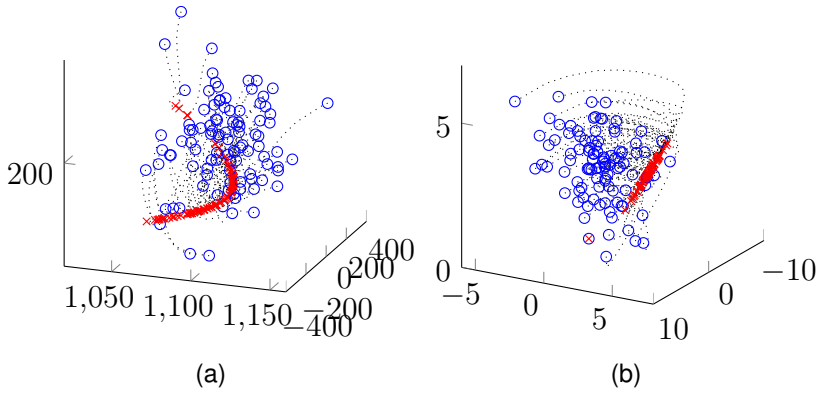

(a) (b)

Figure 11: An example of the GFPF particle motion running on the terrain tracking model (a), showing 3D position, and the skeletal arm model (b), showing 3D shoulder position. Prior states are shown with circles and posterior states with crosses.

and hand positions are calculated using,

$$r_{E,} = r_{S,} + d_U \begin{bmatrix} \cos(\alpha_B)\cos(\alpha_S) \\ \sin(\alpha_S) \\ \sin(\alpha_B)\cos(\alpha_S) \end{bmatrix} \qquad r_{H,} = r_{E,} + d_L \begin{bmatrix} \cos(\alpha_B)\cos(\alpha_S + \alpha_E) \\ \sin(\alpha_S + \alpha_E) \\ \sin(\alpha_B)\cos(\alpha_S + \alpha_E) \end{bmatrix}.$$

Observations of the shoulder and hand positions are made through a perspective projection. By choosing an appropriate coordinate system, this may be modelled simply using,

$$\psi(x) = \begin{bmatrix} \frac{r_{S,1}+r_{S,3}}{r_{S,3}} & \frac{r_{S,2}+r_{S,3}}{r_{S,3}} & \frac{r_{E,1}+r_{E,3}}{r_{E,3}} & \frac{r_{E,2}+r_{E,3}}{r_{E,3}} \end{bmatrix}^T.$$

The observations are accurate, with a variance of 0.001^2.

5.2 Results

Figures 11a and 11b show the motion of the particles from the GFPF on a typical frame, and the awkward shapes of the posterior mode. Tables 1 and 2 show the average ESSs and RMSEs for each algorithm over 100 simulated data sets, each of 100 time steps.

Particle flow resample-move was also tested on the altitude-assisted tracking model. Figure 12 shows the resulting stochastic motion of the particles. Using $\gamma = 0.3$, roughly 25–50% of the MH steps were accepted. The RMSE performance was not significantly improved.

Algorithm	N_F	ESS	RMSE
Bootstrap	5000	1	847
Extended Kalman	1500	40	417
Unscented Kalman	400	18	277
Laplace Approximation	100	14	347
Gaussian Flow	100	57	171

Table 1: Algorithm performance results on the altitude-assisted tracking model, showing number of filter particle (N_F), effective sample size (ESS), and root mean square error (RMSE).

Algorithm	N_F	ESS	RMSE
Bootstrap	11000	1	2.6
Extended Kalman	5000	17	7.2
Laplace Approximation	100	5	6.8
Gaussian Flow	100	58	1.3

Table 2: Algorithm performance results on the skeletal arm model, showing number of filter particle (N_F), effective sample size (ESS), and root mean square error (RMSE). The EPF occasionally diverges and fails to complete. These instances are excluded from the results in the table. The UPF regularly fails and is excluded completely.

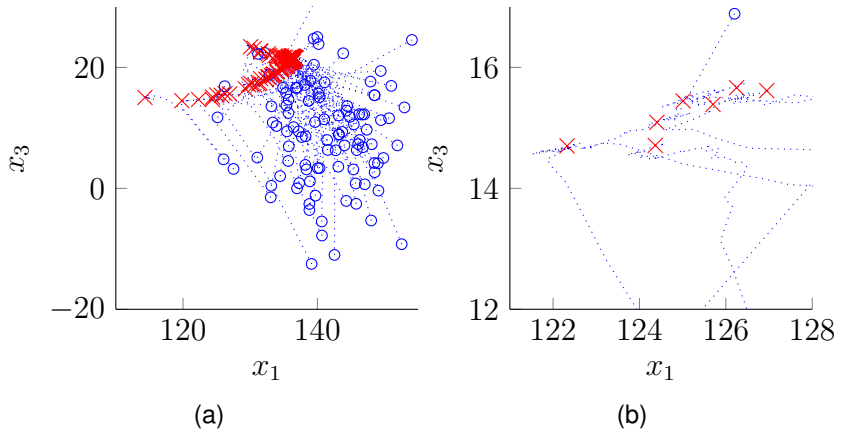

Figure 12: An example of the stochastic GFPF ($\gamma = 0.3$) particle motion running on the altitude-assisted tracking model, showing one horizontal and the vertical state component. Prior states are shown with circles and posterior states with crosses. The second panel is a close-up showing the stochastic motion of the particles.

6 Discussion and Conclusions

We have described the use of particle flow importance sampling using an approximate Gaussian flow, and how this may be used to sample approximately from the optimal importance density of a particle filter. The simulations presented in the previous section demonstrate that this procedure is capable of producing better particle approximations (higher effective sample sizes and lower errors) than simpler particle filters (which use a simple Gaussian importance density) on a class of challenging state space models.

The method introduced is appropriate for models with a Gaussian prior and likelihood but highly nonlinear dependence between the observations and latent state. The algorithm requires almost no tuning. The number of particles and the tolerance for the adaptive step-size selection process are the only critical parameters.

The particle flow and optimal transport methods of (Daum and Huang, 2008, 2011; Reich, 2011, 2012) use similar particle flow ideas to address the task of filtering as we do here. The essential differences in this work are:

- We target the optimal importance density rather than the filtering density directly. The OID is known pointwise up to a normalising constant, and thus we avoid the need for one layer of approximation.
- In (Daum and Huang, 2008, 2011; Reich, 2011, 2012), particle flow samples are used directly to form an approximation of the posterior, with the result that asymptotic convergence properties are lost. We use the particle flow samples as the input to an importance sampler, and correct for the difference between the implied importance density and the posterior density with an appropriate importance weight. Reich (2013) has used a similar importance sampling formulation, but uses different mechanisms to move the particles and assign weights.
- We use an improved numerical integration algorithm based on the analytical solution to the optimal Gaussian flow for linear Gaussian models.

Particle flow algorithms bear a resemblance to annealing-type strategies (Neal, 2001; Deutscher et al., 2000; Gall et al., 2007; Del Moral et al., 2006; Godsill and Clapp, 2001; Oudjane and Musso, 2000), in that both introduce the likelihood progressively. The fundamental difference is that these strategies all use some form of MCMC or resample-move mechanism to remove the weight degeneracy, while particle flow attempts to prevent it happening in the first place. In fact, the two should be seen as complementary. There is no reason why a particle flow could not be used in combination with an annealing scheme. The particles would be moved independently through pseudo-time using a flow, but periodically they are stopped and an intermediate resampling or resample-move step is performed.

Particle flow sampling is only suitable for continuous variables. It should be noted that when the latent state is mixed, with both discrete and continuous components, it is straightforward to sample the discrete component first and then use a particle flow for the continuous part. Furthermore, a number of heavy tailed distributions, including student-t and alpha-stable, can be written as a scale mixture of normals, such that they

are Gaussian conditional on an auxiliary scale variable. If this scale variable is sampled first, then a Gaussian flow may be then be used to sample the state. Successful experiments on such models have been conducted already.

In this work we have exclusively used the methods based on the Gaussian flow, due to its stability and desirable analytical solution. Future research will focus on the use of other choices of particle flow for a more general class of models.

A Particle Flow Governing Equation: Proof of theorem 2.1

The proof follows closely the lines taken by Daum and Huang (2008). First, the log-density is,

$$\Xi_\lambda(x) = M(x) + \lambda L(x) - \log\left(K_\lambda\right),$$

where

$$M(x) = \log\left(p(x)\right) \qquad\qquad L(x) = \log\left(l(x)\right).$$

Differentiating the log of the normalising constant, we find,

$$\begin{aligned}\frac{d}{d\lambda}\log\left(K_\lambda\right) &= \frac{1}{K_\lambda}\frac{dK_\lambda}{d\lambda} \\ &= \frac{\int p(x)l(x)^\lambda L(x)dx}{\int p(x)l(x)^\lambda dx} \\ &= \int \pi_\lambda(x)L(x)dx = \mathbb{E}_{\pi_\lambda}\left[L\right],\end{aligned}$$

and so for the log-density,

$$\frac{\partial \Xi_\lambda}{\partial \lambda} = L(x) - \mathbb{E}_{\pi_\lambda}\left[L\right]. \tag{39}$$

Second, the Fokker-Planck equation relates the motion of a particle with the evolution of the density for its position. For a particle at x_λ moving according to (5) and with density π_λ it states,

$$\begin{aligned}\frac{\partial \pi_\lambda}{\partial \lambda} &= -\operatorname{Tr}\left[\frac{\partial}{\partial x_\lambda}\left(\zeta_\lambda(x_\lambda)\pi_\lambda(x_\lambda)\right)\right] + \sum_{ij}\frac{\partial^2}{\partial x_{\lambda,i}\partial x_{\lambda,j}}\left(D_{\lambda,ij}(x_\lambda)\pi_\lambda(x_\lambda)\right) \\ &= -\pi_\lambda(x_\lambda)\operatorname{Tr}\left[\frac{\partial \zeta_\lambda}{\partial x_\lambda}\right] - \frac{\partial \pi_\lambda}{\partial x_\lambda}^T\zeta_\lambda(x_\lambda) + \operatorname{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \pi_\lambda}{\partial x_\lambda^2}\right] \\ &\qquad + 2\sum_{ij}\frac{\partial D_{\lambda,ij}}{\partial x_{\lambda,i}}\frac{\partial \pi_\lambda}{\partial x_{\lambda,j}} + \pi_\lambda(x_\lambda)\sum_{ij}\frac{\partial^2 D_{\lambda,ij}}{\partial x_{\lambda,i}\partial x_{\lambda,j}},\end{aligned} \tag{40}$$

where $D_\lambda(x) = \frac{1}{2}\eta_\lambda(x)\eta_\lambda(x)^T$. This may be recast to use log-densities using the following identities,

$$\frac{\partial \Xi_\lambda}{\partial \lambda} = \frac{1}{\pi_\lambda(x)}\frac{\partial \pi_\lambda}{\partial \lambda} \qquad \frac{\partial \Xi_\lambda}{\partial x} = \frac{1}{\pi_\lambda(x)}\frac{\partial \pi_\lambda}{\partial x}$$

$$\frac{\partial^2 \Xi_\lambda}{\partial x^2} = \frac{\pi_\lambda(x)\frac{\partial^2 \pi_\lambda}{\partial x^2} - \frac{\partial \pi_\lambda}{\partial x}\frac{\partial \pi_\lambda}{\partial x}^T}{\pi_\lambda(x)^2} = \frac{1}{\pi_\lambda(x)}\frac{\partial^2 \pi_\lambda}{\partial x^2} - \frac{\partial \Xi_\lambda}{\partial x}\frac{\partial \Xi_\lambda}{\partial x}^T.$$

Dividing (40) through by $\pi_\lambda(x_\lambda)$ (assuming that this is nowhere vanishing) we obtain,

$$\begin{aligned}\frac{\partial \Xi_\lambda}{\partial \lambda} = &- \operatorname{Tr}\left[\frac{\partial \zeta_\lambda}{\partial x_\lambda}\right] - \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) + \operatorname{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \Xi_\lambda}{\partial x_\lambda^2}\right] + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T D_\lambda(x_\lambda)\frac{\partial \Xi_\lambda}{\partial x_\lambda} \\ &+ 2\sum_{ij}\frac{\partial D_{\lambda,ij}}{\partial x_{\lambda,i}}\frac{\partial \Xi_\lambda}{\partial x_{\lambda,j}} + \sum_{ij}\frac{\partial^2 D_{\lambda,ij}}{\partial x_{\lambda,i}\partial x_{\lambda,j}}. \end{aligned} \tag{41}$$

Combining the equations for the log-density (39) with the partial differential equation for the log-density evolution (41), the governing equation for the optimal particle dynamics is reached. ∎

B Evolution of Ideal Importance Weights: Proof of theorem 2.2

Define $\Upsilon_\lambda(x_\lambda) = \log(\upsilon_\lambda(x_\lambda))$, and apply Itō's lemma,

$$d\Upsilon_\lambda = \left[\frac{\partial \Upsilon_\lambda}{\partial \lambda} + \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) + \frac{1}{2}\operatorname{Tr}\left[\eta_\lambda \eta_\lambda^T \frac{\partial^2 \Upsilon_\lambda}{\partial x_\lambda^2}\right]\right] d\lambda + \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}^T \eta_\lambda d\epsilon_\lambda.$$

Equivalently to (41), the Fokker-Planck equation tells us that,

$$\begin{aligned}\frac{\partial \Upsilon_\lambda}{\partial \lambda} = &- \operatorname{Tr}\left[\frac{\partial \zeta_\lambda}{\partial x_\lambda}\right] - \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) + \operatorname{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \Upsilon_\lambda}{\partial x_\lambda^2}\right] + \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}^T D_\lambda(x_\lambda)\frac{\partial \Upsilon_\lambda}{\partial x_\lambda} \\ &+ 2\sum_{ij}\frac{\partial D_{\lambda,ij}}{\partial x_{\lambda,i}}\frac{\partial \Upsilon_\lambda}{\partial x_{\lambda,j}} + \sum_{ij}\frac{\partial^2 D_{\lambda,ij}}{\partial x_{\lambda,i}\partial x_{\lambda,j}}. \end{aligned}$$

Combining these,

$$d\Upsilon_\lambda = \Bigg[-\operatorname{Tr}\left[\frac{\partial \zeta_\lambda}{\partial x_\lambda}\right] + 2\ \operatorname{Tr}\left[D_\lambda(x_\lambda)\frac{\partial^2 \Upsilon_\lambda}{\partial x_\lambda^2}\right] + \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}^T D_\lambda(x_\lambda)\frac{\partial \Upsilon_\lambda}{\partial x_\lambda} + 2\sum_{ij}\frac{\partial D_{\lambda,ij}}{\partial x_{\lambda,i}}\frac{\partial \Upsilon_\lambda}{\partial x_{\lambda,j}} + \sum_{ij}\frac{\partial^2 D_{\lambda,ij}}{\partial x_{\lambda,i}\partial x_{\lambda,j}}\Bigg] d\lambda + \frac{\partial \Upsilon_\lambda}{\partial x_\lambda}^T \eta_\lambda d\epsilon_\lambda.$$

Next, using Itō'e Lemma for the target sequence log-density, and inserting (39),

$$\begin{aligned} d\Xi_\lambda &= \left[\frac{\partial \Xi_\lambda}{\partial \lambda} + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) + \frac{1}{2}\operatorname{Tr}\left[\eta_\lambda \eta_\lambda^T \frac{\partial^2 \Xi_\lambda}{\partial x_\lambda^2}\right]\right] d\lambda + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \eta_\lambda d\epsilon_\lambda \\ &= \left[L(x_\lambda) - \mathbb{E}_{\pi_\lambda}[L] + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \zeta_\lambda(x_\lambda) + \operatorname{Tr}\left[D_\lambda \frac{\partial^2 \Xi_\lambda}{\partial x_\lambda^2}\right]\right] d\lambda + \frac{\partial \Xi_\lambda}{\partial x_\lambda}^T \eta_\lambda d\epsilon_\lambda. \end{aligned} \quad (42)$$

Finally, for the log-weight $W_\lambda = \log(w_\lambda)$, we have,

$$dW_\lambda = d\Xi_\lambda - d\Upsilon_\lambda,$$

and substituting the two differentials the results is reached. ∎

C Evolution of Practical Importance Weights: Proof of theorem 2.3

It is well known that any diffusion process may be constructed as the limit of a particular discrete time Markov chain as the step size tends to 0 (Øksendal, 2003). Specifically, for an Itō diffusion, if we have time instants at $\lambda_n = n\delta\lambda$, then,

$$x_n = x_{n-1} + \zeta_{n-1}(x_{n-1})\delta\lambda + \eta_{n-1}(x_{n-1}) z_n \delta\lambda^{\frac{1}{2}}, \quad (43)$$

where $\{z_n\}$ are drawn independently from a standard Gaussian distribution with density $\varsigma(z) = \mathcal{N}(z\,|0,\ \delta\lambda I)$. We can derive an appropriate differential equation for a particle importance weight by constructing a sequential importance sampler on this system and then taking the limit $\delta\lambda \to 0$.

For a particle with density υ_λ to be properly weighted with respect to the target density π_λ, the ideal importance weight is given in (8). This is not practical because υ_λ is generally intractable. Instead we construct an extended target distribution over $\{x_n, z_1, z_2, \ldots, z_n\}$, in the manner of a sequential Monte Carlo (SMC) sampler (Del Moral et al., 2006),

$$\pi_n(x_n)\prod_{k=1}^{n}\varsigma(z_k). \quad (44)$$

Samples are drawn by first simulating x_0 from p, the state prior, and z_k for $k = 1, \ldots, n$ from ς, and then recursively applying (43). We can write the inverse of this transformation using Taylor series expansions of ζ_{n-1} and η_{n-1},

$$x_{n-1} = x_n - \zeta_n(x_n)\delta\lambda - \eta_n(x_n)z_n\delta\lambda^{\frac{1}{2}} + \frac{\partial\left[\eta_n(x_n)z_n\right]}{\partial x_n}\eta_n(x_n)z_n\delta\lambda + \mathcal{O}\left(\delta\lambda^{\frac{3}{2}}\right).$$

By the change of variables formula, and using the Jacobian of this inverse transformation, the proposal density in the extended space is,

$$p(x_0)\prod_{k=1}^{n}\varsigma(z_k)\left|\frac{\partial x_{k-1}}{\partial x_k}\right|. \tag{45}$$

The resulting importance weight is the ratio of target (44) and proposal (45) densities in the extended space. Taking the log,

$$\begin{aligned} W_n &= \Xi_n(x_n) - M(x_0) - \sum_{k=1}^{n}\log\left(\left|\frac{\partial x_{k-1}}{\partial x_k}\right|\right) \\ &= W_{n-1} + \Xi_n(x_n) - \Xi_{n-1}(x_{n-1}) - \log\left(\left|\frac{\partial x_{n-1}}{\partial x_n}\right|\right). \end{aligned} \tag{46}$$

To calculate the Jacobian term, we use the following identities,

$$\begin{aligned} |I + \delta A| &= 1 + \delta\ \mathrm{Tr}\left[A\right] + \frac{1}{2}\delta^2\left(\mathrm{Tr}\left[A\right]^2 - \mathrm{Tr}\left[A^2\right]\right) + \mathcal{O}\left(\delta^3\right) \\ \log(1+\delta a) &= \delta a - \frac{\delta^2}{2}a^2 + \mathcal{O}\left(\delta^3\right), \end{aligned}$$

with which we reach,

$$-\log\left(\left|\frac{\partial x_{n-1}}{\partial x_n}\right|\right) = \mathrm{Tr}\left[\frac{\partial\zeta_{n-1}}{\partial x_{n-1}}\right]\delta\lambda + \sum_{ij}\frac{\partial\eta_{n-1,ij}}{\partial x_{n-1,i}}\delta\epsilon_{n,j} - \frac{1}{2}\sum_{ijk}\left[\frac{\partial\eta_{n-1,ik}}{\partial x_{n-1,j}}\frac{\partial\eta_{n-1,jk}}{\partial x_{n-1,i}}\right]\delta\lambda + \mathcal{O}\left(\delta\lambda^{\frac{3}{2}}\right).$$

Finally, letting $\delta\lambda \to 0$ and substituting (42) into (46), we obtain the result. ∎

D Integrated Gaussian Flow: Proof of Theorem 3.4

For a small increment of pseudo-time, such that $\lambda_0 = \lambda$ and $\lambda_1 = \lambda + \delta\lambda$,

$$x_{\lambda+\delta\lambda} = m_{\lambda+\delta\lambda} + \exp\left\{-\frac{1}{2}\gamma\delta\lambda\right\}\left(P_{\lambda+\delta\lambda}P_\lambda^{-1}\right)^{\frac{1}{2}}\left(x_\lambda - m_\lambda\right) + \left[\frac{1-\exp\left\{-\gamma\delta\lambda\right\}}{\delta\lambda}\right]^{\frac{1}{2}}P_{\lambda+\delta\lambda}^{\frac{1}{2}}\left(\epsilon_{\lambda_1} - \epsilon_{\lambda_0}\right).$$

Now use the following expansions for small increments,

$$\begin{aligned}
m_{\lambda+\delta\lambda} &= m_\lambda + \frac{\partial m_\lambda}{\partial \lambda}\delta\lambda + \mathcal{O}\left(\delta\lambda^2\right) \\
P_{\lambda+\delta\lambda}^{\frac{1}{2}} &= P_\lambda^{\frac{1}{2}} + \frac{\partial P_\lambda^{\frac{1}{2}}}{\partial \lambda}\delta\lambda + \mathcal{O}\left(\delta\lambda^2\right) \\
P_{\lambda+\delta\lambda} &= \left(P_\lambda^{-1} + \delta\lambda H^T R^{-1} H + \mathcal{O}\left(\delta\lambda^2\right)\right)^{-1} \\
&= P_\lambda - \delta\lambda P_\lambda H^T R^{-1} H P_\lambda + \mathcal{O}\left(\delta\lambda^2\right) \\
\left(P_{\lambda+\delta\lambda} P_\lambda^{-1}\right)^{\frac{1}{2}} &= \left(I - \delta\lambda P_\lambda H^T R^{-1} H + \mathcal{O}\left(\delta\lambda^2\right)\right)^{\frac{1}{2}} \\
&= I - \frac{1}{2}\delta\lambda P_\lambda H^T R^{-1} H + \mathcal{O}\left(\delta\lambda^2\right) \\
\exp\left\{-\frac{1}{2}\gamma\delta\lambda\right\} &= 1 - \frac{1}{2}\gamma\delta\lambda + \mathcal{O}\left(\delta\lambda^2\right) \\
\left[\frac{1-\exp\left\{-\gamma\delta\lambda\right\}}{\delta\lambda}\right]^{\frac{1}{2}} &= \left[\frac{\gamma\delta\lambda + \mathcal{O}\left(\delta\lambda^2\right)}{\delta\lambda}\right]^{\frac{1}{2}} = \gamma^{\frac{1}{2}} + \mathcal{O}\left(\delta\lambda\right),
\end{aligned}$$

and noting that,

$$\frac{\partial m_\lambda}{\partial \lambda} = P_\lambda H^T R^{-1}\left(y - H m_\lambda\right),$$

leads to,

$$x_{\lambda+\delta\lambda} - x_\lambda = P_\lambda H^T R^{-1}\left[(y - H m_\lambda) - \frac{1}{2}H(x_\lambda - m_\lambda) - \frac{1}{2}\gamma(x_\lambda - m_\lambda)\right]\delta\lambda + \gamma^{\frac{1}{2}} P_\lambda^{\frac{1}{2}}\left(\epsilon_{\lambda+\delta\lambda} - \epsilon_\lambda\right) + \mathcal{O}\left(\delta\lambda\right)$$

Taking the limit as $\delta\lambda \to 0$, the result follows. ∎

E Weight Numerical Integration: Proof of Theorem 3.6

Using Taylor expansions,

$$\log\left(\left|\frac{\partial x_{\lambda_1}}{\partial x_{\lambda_0}}\right|\right) = -\log\left(\left|\frac{\partial x_{\lambda_0}}{\partial x_{\lambda_1}}\right|\right) + \mathcal{O}\left(\left(\lambda_1 - \lambda_0\right)^{\frac{3}{2}}\right).$$

Hence, the log-weight update for a small step is described by (46) with an error of order $\mathcal{O}\left(\left(\lambda_1 - \lambda_0\right)^{\frac{3}{2}}\right)$. Therefore, by construction, as the step size goes to 0, the weight evolves according to (10), and from theorem 2.3, the particle is properly weighted. ∎

References

Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72:269–342.

Bar-Shalom, Y. and Li, X. R. (1995). *Multitarget-multisensor tracking: principles and techniques*. Storrs, CT : Yaakov Bar-Shalom.

Bartels, R. H. and Stewart, G. W. (1972). Solution of the matrix equation ax + xb = c. *Commun. ACM*, 15(9):820–826.

Bunch, P. and Godsill, S. (2013). Particle filtering with progressive Gaussian approximations to the optimal importance density. In *5th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*.

Cappé, O., Godsill, S., and Moulines, E. (2007). An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE*, 95:899–924.

Daum, F. and Huang, J. (2008). Particle flow for nonlinear filters with log-homotopy. In *Proceedings of SPIE, the International Society for Optical Engineering*. Society of Photo-Optical Instrumentation Engineers.

Daum, F. and Huang, J. (2011). Particle degeneracy: root cause and solution. In *Proc. SPIE*, volume 8050. SPIE.

Daum, F. and Huang, J. (2013). Particle flow with non-zero diffusion for nonlinear filters. In *Proceedings of SPIE 8745, Signal Processing, Sensor Fusion, and Target Recognition XXII*, volume 8745. SPIE.

Daum, F., Huang, J., Krichman, M., and Kohen, T. (2009). Seventeen dubious methods to approximate the gradient for nonlinear filters with particle flow. In Drummond, O. E. and Teichgraeber, R. D., editors, *Proc. SPIE*, volume 7445. SPIE.

Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436.

Deutscher, J., Blake, A., and Reid, I. (2000). Articulated body motion capture by annealed particle filtering. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 126–133.

Douc, R., Cappe, O., and Moulines, E. (2005). Comparison of resampling schemes for particle filtering. In *Proc. 4th Int. Symp. Image and Signal Processing and Analysis*.

Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10:197–208.

Doucet, A. and Johansen, A. M. (2009). A tutorial on particle filtering and smoothing: Fifteen years later. In Crisan, D. and Rozovsky, B., editors, *The Oxford Handbook of Nonlinear Filtering*. Oxford University Press.

Gall, J., Potthoff, J., Schnörr, C., Rosenhahn, B., and Seidel, H.-P. (2007). Interacting and annealing particle filters: Mathematics and a recipe for applications. *Journal of Mathematical Imaging and Vision*, 28(1):1–18.

Gelman, A. and Meng, X.-L. (1998). Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163–185.

Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, 57(6):1317–1339.

Gilks, W. R. and Berzuini, C. (2001). Following a moving target — Monte Carlo inference for dynamic Bayesian models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(1):127–146.

Godsill, S. and Clapp, T. (2001). Improvement strategies for Monte Carlo particle filters. In Doucet, A., de Freitas, N., and Gordon, N., editors, *Sequential Monte Carlo Methods in Practice*, pages 139–158. Springer New York.

Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F, Radar and Signal Processing*, 140(2):107–113.

Hagmar, J., Jirstrand, M., Svensson, L., and Morelande, M. (2011). Optimal parameterization of posterior densities using homotopy. In *14th International Conference on Information Fusion (FUSION)*.

Hanebeck, U. and Steinbring, J. (2012). Progressive Gaussian filtering. *arXiv preprint arXiv:1204.0133*.

Hanebeck, U. D. and Feiermann, O. (2003). Progressive Bayesian estimation for nonlinear discrete-time systems:the filter step for scalar measurements and multidimensional states. In *42nd IEEE Conference on Decision and Control*, volume 5, pages 5366–5371.

Hol, J. D., Schon, T. B., and Gustafsson, F. (2006). On resampling algorithms for particle filters. In *IEEE Nonlinear Statistical Signal Processing Workshop*, pages 79–82.

Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25.

Kong, A., Liu, J. S., and Wong, W. H. (1994). Sequential imputations and bayesian missing data problems. *Journal of the American statistical association*, 89(425):278–288.

Liu, J. S. (2001). *Monte Carlo strategies in scientific computing*. Springer.

Neal, R. M. (2001). Annealed importance sampling. *Statistics and Computing*, 11(2):125–139.

Øksendal, B. (2003). *Stochastic differential equations: An introduction with applications*. Springer Verlag.

Oudjane, N. and Musso, C. (2000). Progressive correction for regularized particle filters. In *3rd International Conference on Information Fusion (FUSION)*, volume 2, pages 10–17.

Reich, S. (2011). A dynamical systems framework for intermittent data assimilation. *BIT Numerical Mathematics*, 51:235–249.

Reich, S. (2012). A Gaussian-mixture ensemble transform filter. *Quarterly Journal of the Royal Meteorological Society*, 138(662):222–233.

Reich, S. (2013). A guided sequential monte carlo method for the assimilation of data into stochastic dynamical systems. In *Recent Trends in Dynamical Systems*, pages 205–220. Springer.

Schön, T., Gustafsson, F., and Nordlund, P.-J. (2005). Marginalized particle filters for mixed linear/nonlinear state-space models. *IEEE Transactions on Signal Processing*, 53(7):2279–2289.

Shampine, L. and Reichelt, M. (1997). The MATLAB ODE suite. *SIAM Journal on Scientific Computing*, 18(1):1–22.

Van Der Merwe, R., Doucet, A., Freitas, N., and Wan, E. (2000). The unscented particle filter. *Advances in Neural Information Processing Systems*, 13.